

# Mechanism Design for Value Maximizers<sup>1</sup>

CHRISTOPHER A WILKENS, Yahoo Research  
 RUGGIERO CAVALLO, Yahoo Research  
 RAD NIAZADEH, Cornell University

Bidders often want to get as much as they can without violating constraints on what they spend. For example, advertisers seek to maximize the impressions, clicks, sales, or market share generated by their advertising, subject to budget or return-on-investment (ROI) constraints. Likewise, when bidders have no direct utility for leftover money – e.g., because the money comes from a corporate budget – they will naturally buy as much as possible. We call such bidders *value maximizers*. The quasilinear utility model dramatically fails to capture these preferences, and so we initiate the study of mechanism design in this different context. In single-parameter settings, we show that any monotone allocation can be implemented truthfully. Interestingly, even in unrestricted domains any choice function that optimizes monotone functions of bidders' values can be implemented; this contrasts with the case of quasilinear preferences, where Roberts [1979] showed that only affine maximizers are implementable. For general valuations we show that maximizing the value of the highest-value bidder (and iterating through successively lower-value bidders to break ties) is the natural analog of welfare maximization.

We apply our results to online advertising as a case study. Firstly, the natural analog of welfare maximization directly generalizes the generalized second price (GSP) auction commonly used to sell search ads. We advocate for this generalization, discussing how it and related hybrid auctions can be used to properly handle the complex ad options available in sponsored search today. Finally, we show that value-maximizing preferences are robust in a practical sense: even though real advertisers' preferences are complex and varied, as long as outcomes are “sufficiently different,” any advertiser with a moderate ROI constraint and “super-quasilinear” preferences will be behaviorally equivalent to a value maximizer. We empirically evaluate the implications of this, establishing that for at least 80% of a sample of auctions from the Yahoo Gemini Native ads platform, bidders requiring at least a 100% ROI should behave like value maximizers.

## 1. INTRODUCTION

Online advertising is a major success story for mechanism design. Starting in the mid-1990s, enterprising websites developed novel mechanisms for selling ads; in the 2000s, academics across many disciplines recognized them as auctions and projected their models onto this exciting new world, designing exquisite mechanisms and offering deep insights into existing ones. Yet, at least one issue seems to have escaped scrutiny: the research community rationalized its standard model of quasilinear utilities — which capture profit in their simplest form — as the natural incarnation of advertisers' objectives. It is widely acknowledged that advertisers may have complex preferences and behaviors (e.g. [Aggarwal et al. 2009]), but it has gone unquestioned that the quasilinear model is the natural place to *start* designing mechanisms. But from a vantage point within a company that sells advertising, this strains credulity;

<sup>1</sup>A preliminary work “On the Truthfulness of GSP” [Cavallo et al. 2015] was presented at the 11th Ad Auctions Workshop.

---

Authors' email addresses: cwilkens@yahoo-inc.com, cavallo@yahoo-inc.com, rad@cs.cornell.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'16, July 24–28, 2016, Maastricht, The Netherlands. Copyright © 2016 ACM 978-1-4503-3936-0/16/07 ...\$15.00.

considering requests from both advertisers and the account managers who advocate on their behalf, the way buyers interact with a marketplace and make local bidding decisions is very far from profit maximization. Rather, we suggest that advertiser behavior is more fundamentally driven by a desire to maximize the value of what they buy — total impressions, clicks, sales, market share, etc. — subject to budget constraints and return on investment (ROI) demands. Moreover, advertisers are not unique; this behavior arises naturally in many common situations where budgets are fixed *ex ante* or profit is difficult to measure.

We propose that bidders such as advertisers should be formally modeled in single-shot mechanisms as *value maximizers*. Whereas the standard quasilinear model says that a bidder has a value  $v_i$  and strives to maximize value minus price  $v_i - p_i$ , a value maximizer instead strives simply to maximize  $v_i$ , as long as  $p_i$  satisfies bidder specific constraints, though a lower price is always preferred for the same value. One natural assumption when  $v_i$  is measured in dollars is that bidders always want to spend less than they make, i.e., they always require  $p_i \leq v_i$ . This gives us our definition of a simple value maximizer:

*Definition 1.1 (Simple Value-Maximizer, Informal).* A simple value maximizer tries to maximize value  $v_i$  while keeping the total payment  $p_i$  less than its value, i.e., to achieve an outcome that maximizes  $v_i$  while keeping  $p_i \leq v_i$ .

In practice, any value maximizer will have constraints in addition to  $p_i \leq v_i$ . Two types are of paramount importance here: ROI constraints and budget constraints. Budget constraints require that  $p_i \leq B_i$ . ROI measures the ratio of the profit obtained (“return”) to the price paid (“investment”), i.e., the density of profit in cost:

$$ROI = \frac{Value - Price}{Price} = \frac{Profit}{Price}$$

An ROI constraint requires that this ratio meet some minimum value  $\gamma$ , i.e.,  $\frac{v_i - p_i}{p_i} \geq \gamma$ . ROI is a common metric in business and, for both theoretical and practical reasons, ROI constraints are the natural place to start our study:

- *ROI constraints are standard in advertising.* As we will discuss, an advertiser who measures a cost-benefit trade-off typically uses ROI as the metric to do it.
- *ROI constraints are a heuristic for implementing budget constraints.* When each auction has a limited impact on an agent’s budget, ROI constraints are a common heuristic for setting bids across auctions, particularly in advertising [Borgs et al. 2007; Kitts and Leblanc 2004; Szymanski and Lee 2006; Zhou et al. 2008; Auerbach et al. 2008]. Thus, since our focus is on single-shot auctions, ROI constraints are the natural place to start.
- *ROI constraints are theoretically clean.* Return on investment (ROI) is a slight variant on the simple value maximizer: a value maximizer who requires ROI at least  $\gamma$  would require  $p_i \leq \frac{v_i}{1+\gamma}$ .<sup>2</sup> Since a bidder with an ROI constraint is like a bidder with a reduced value  $v'_i = \frac{v_i}{1+\gamma}$ , many results from the simple setting will apply directly to ROI-constrained bidders. In contrast, it is well-known that budgets dramatically complicate mechanism design in quasilinear settings, and value-maximization settings will not be an exception.

We study a range of mechanism design problems from fundamental theory to practical ad auctions. First, we ask basic questions about the existence of truthful mechanisms and extend standard characterizations to simple value maximizers. Next, we

<sup>2</sup>ROI is computed as  $\frac{v_i - p_i}{p_i}$ , so  $ROI \geq \gamma$  implies  $p_i \leq \frac{v_i}{1+\gamma}$ .

apply those basic results to online advertising auctions as a case study. We gain insight into the GSP auction commonly used in practice and develop new auctions that handle the growing complexity in search ads. Third, we take a decidedly practical approach, and study mechanism design for bidders who do not precisely fit any of our existing moulds. We show that value-maximizing preferences are robust in the sense that any bidder with an ROI constraint will look like a value maximizer under certain circumstances; we show that these circumstances are not a theoretical curiosity and that they plausibly exist for a majority of auctions in a sample from Yahoo’s Gemini Native advertising platform. Our work is by no means exhaustive and leaves many thorny problems, particularly related to agent preferences for lotteries and across auctions; we conclude with significant open questions it raises.

### **Advertiser Objectives in Practice**

Advertisers have many objectives in practice, but we will argue that most point towards value-maximizing preferences as the natural starting point for mechanism design.

One class of advertisers target growth and reach; they are intuitively modeled as value maximizers. These advertisers want to reach as many people as they can, generate as many clicks as they can, capture as much market share as they can, and so on. These objectives are often designed to meet business objectives other than profit (e.g. revenue growth), are difficult to map to money (e.g. market share), or are difficult to directly measure (e.g. monetary impact of display advertising). Thus, instead of managing a delicate trade-off between value and money, it is natural that such an advertiser tries to purchase as much advertising as possible without violating coarse spending constraints.

When advertisers do consider cost-benefit trade-offs as a primary concern, ROI is the standard metric across all types of advertising. From the outside, profit maximization seems like a reasonable description of what an advertiser should want; however, hearing from advertisers and their representatives, it is clear that their behavior and (at least short-term) objectives are typically not focused on profit. For example, the following story plays out regularly at Yahoo:

- (1) Advertiser X designates a small budget for testing a Yahoo advertising product.
- (2) Advertiser X measures ROI — if X is happy with its ROI, it increases its budget hoping to *maintain the same ROI*; if it is unhappy, it withdraws from the marketplace.

This behavior is common across all sizes and types of performance-driven advertiser, and it is consistent with ROI as the driving factor over and above profit. The standard industry tools also paint a similar picture: Google’s AdWords campaign management tool buys as much advertising as possible while maintaining a target average CPC and budget — an average CPC constraint is appropriate for ROI-centric preferences, while a marginal CPC constraint would be appropriate for profit-centric preferences. Yahoo’s display ad products also offer options with “CPC goals” and “CPA goals” that buy impressions while aiming to achieve a target average cost-per-click or cost-per-conversion (e.g., cost per purchase) respectively.

All of these advertisers face a problem of budget management, and, again, budget management tends to manifest as ROI constraint at the auction level. When faced with a problem of distributing budget among individual auctions, ROI is a natural and common basis for building heuristics [Borgs et al. 2007; Kitts and Leblanc 2004; Szymanski and Lee 2006; Zhou et al. 2008; Auerbach et al. 2008]. Thus, when we model advertiser incentives in a *single auction*, budgets point towards bidders with ROI constraints.

We formalize ROI-centric preferences and find, perhaps surprisingly, that they naturally map to a value-maximizing model. Recall that ROI is the density of profit in cost:  $ROI = \frac{Value - Price}{Price} = \frac{Profit}{Price}$ . An immediate observation is that ROI alone is insufficient to characterize performance, since ROI can be high even when an advertiser is buying almost nothing. Instead, advertisers must be modeled as solving a constrained optimization problem, where ROI is combined with total value or total cost. Note the contrast with profit-centric preferences — maximizing profit is sensible even without added constraints.

With ROI in mind, two natural models stand out when we substitute ROI for profit:

- *ROI as a constraint; value as an objective.* First and foremost, advertisers use ROI as a constraint. An advertiser with a ROI constraint naturally tries to buy as much advertising as possible subject to a constraint on ROI (and likely budget).
- *ROI as an objective; value or cost as a constraint.* An advertiser who makes ROI an objective must fix the amount of money they wish to spend or the total value they wish to get. The former is more common and is equivalent to maximizing the total amount of value subject to a fixed amount of money spent.

Both of these scenarios point towards the same fundamental model of preferences: advertisers buy as much advertising as possible subject to constraints on ROI and total value.

Thus, we reach the conclusion that *advertisers are naturally modeled as bidders who prefer more value as long as ROI, budget, and value constraints are satisfied*. There are always exceptions; anecdotal evidence suggests that a savvy minority of small businesses do target profit. In small businesses, advertising is often managed by the business owner, so it is easier to directly tie advertising to profit. Likewise, there are always unusual objectives, such as an advertiser who (rationally or irrationally) always wants the top slot on Google. Such examples will exist for any model; we contend that value maximization is simple and worthy of dedicated study.

### **Value Maximizers Beyond Advertising**

Value maximization is an extremely common behavior; it is not particular to advertising. Value maximization arises naturally any time agents spend money that is not their own. This is particularly common in large organizations like companies and governments — an employee is given a budget for a particular task (say, to provide food for an event), and the employee is typically rewarded for doing that task well, not for saving money. Since the money is not the employee's, she won't get any benefit from underspending the budget. Thus, the only incentive is to maximize value of what is provided (say, quality or quantity of food) with the budget as a constraint. In a sense, advertising is simply a special case of this behavior.

Value maximization also arises naturally when long-term value is unclear. While a company would like to know the precise return that an investment will generate — such as acquiring another company, advertising, or launching a new product — this is often nearly impossible to know. As a result, it is not possible to make the precise value-cost tradeoff required to optimize profit, and it is safer to get the most value from an investment while keeping costs within “sensible” bounds. Again, advertising is simply a special case.

Individuals also maximize value with their own money. People often fix a budget before making a purchase, then find the best item that meets the budget. Buying a house is an easy example: buyers will commonly fix a budget and then push that budget to get as much “value” out of the house as possible — number of rooms, land, upgrades, etc. — without violating the budget.

### **Mechanism Design for Value Maximizers**

Since value maximizers do not fit the standard models in the current literature, we first revisit foundational results in mechanism design, and extend the standard single-parameter and affine-maximizer characterizations to value maximizers. An immediate challenge is that value maximizers have a nearly lexicographic preference that cannot be represented by a utility function — they prefer value first, and only care about price for a fixed value.

For single-parameter domains, we show that an allocation rule can be implemented if and only if it is monotone, just as in the standard quasilinear setting [Myerson 1981; Archer and Tardos 2001]. The price a bidder pays is the minimum bid required to receive the same value. Our proof is a general characterization that applies to a large family of bidder preferences, not just value maximizers.

For general domains, we show that an analog of the celebrated Vickrey-Clarke-Groves mechanism (VCG) always exists. In quasilinear settings, maximizing the  $L^1$ -norm of agents' values is known as maximizing welfare; it is always incentive compatible. For value maximizers, we show that the analogous mechanism loosely maximizes the  $L^\infty$  norm of agents' values — it maximizes the value of the highest-value bidder, then the value of the second-highest-value bidder, and so on.

Interestingly, we show that more mechanisms are implementable in unrestricted domains than in quasilinear settings. In quasilinear settings, Roberts's theorem [Roberts 1979] says that affine variants of welfare are the only objectives that can be maximized truthfully; interestingly, we show that truthful mechanisms we can maximize any bidder-specific monotone transformations of value. We demonstrate that this is a fragile feature of value maximizers; for a family of utility functions for which maximizing the  $L^p$ -norm is truthful (which loosely converges to value maximizers as  $p \rightarrow \infty$ ), we show that an analog of Roberts's theorem still holds.

Many of these results come with a technical caveat about ties — if a tie would occur in a continuous bid space, then agents will have an incentive to lie. This is an unavoidable consequence of the particular lexicographic form of bidders' preferences but happens very rarely, so we accept that those mechanisms will only be truthful “almost everywhere” over bidders' typespaces.

### Mechanism Design for Advertising

As a case study, we apply our fundamental results to online advertising, and learn key insights about the generalized second price auction (GSP) used in practice that help solve pressing problems in search advertising.

We first observe an ironic twist of fate: the generalized second price auction (GSP) — used in practice but derided as non-truthful in theory — is the incentive compatible auction for value maximizers. This observation was first made by Aggarwal et al. [Aggarwal et al. 2009] who studied matching mechanisms (as a generalization of sponsored search auctions) under a class of valuations that included value maximizers. In our context, we observe that GSP is the special case of the “welfare-maximizing” auction for the standard sponsored search model.

Moreover, we take this observation a step farther and suggest that GSP should be defined this way — *GSP is the auction that is truthful for value maximizers with an ROI constraint.*

Generalizing GSP in this way helps us solve a problem facing search engines today — what should be done with complex ad types? Search ads today often come with many different features and in many different sizes, something that cannot be handled properly by a standard GSP auction. Our definition tells us what a GSP auction should do in these complex settings and how prices should be computed. Moreover, if this version of GSP is unsatisfactory, it suggests a natural way to interpolate between GSP and VCG: construct a family of preferences that interpolate between quasilinear and

value-maximizing, then run a truthful auction for a set of preferences in that family. We use the family of preferences used in our characterization theorems and discuss how they might be used for this purpose.

### Mechanism Design for Real Advertisers

Finally, we discover that value-maximizing preferences are robust in a useful sense: *ROI constraints make almost any advertiser look like a value maximizer* under plausible constraints. We show that in any setting, as long as bidders’ values are sufficiently different, We study the standard GSP auction for advertisers who have “super-quasilinear” preferences (roughly, anything between profit-maximizing and value-maximizing preferences) and show that adding a sufficient ROI constraint makes that advertiser behave like a pure value maximizer as long as the available bundles of goods are sufficiently differentiated. Moreover, this is not merely a theoretical curiosity: based on data from Yahoo’s Gemini Native advertising platform, we show that 80% of auctions would be truthful under GSP pricing given only that advertisers required an  $ROI > 1$ . This suggests that it may be reasonable in practice to assume that bidders are value maximizers.

### Related Work

There is a limited relevant literature that studies mechanism design for bidders with general non-quasilinear utilities. The most closely related work to our own is [Aggarwal et al. 2009]. They design sponsored search (slot) auctions for advertisers whose preferences add constraints on top of a quasilinear utility. While their general models are technically incomparable to ours, they do cover our value maximizers in the special case of sponsored search, and they make the observation that the truthful auction is simply GSP. Another work that relaxes the quasilinearity assumption for ad auctions is [Alaei et al. 2011]. The main focus of their work is on Walrasian equilibria in unit-demand settings, but they show an application to ad auctions where estimation errors break quasilinearity for utilities that would otherwise be quasilinear.

Outside of advertising, a few papers study general truthful mechanisms for non-quasilinear preferences [Adachi 2013; Morimoto and Serizawa 2015]. These papers develop axiomatic characterizations of a VCG-analog for multi-unit and unit-demand settings.

There is also a literature that considers heuristics of advertiser behavior. One line of work studies the effect of using ROI to set bids across different auctions [Borgs et al. 2007; Kitts and Leblanc 2004; Szymanski and Lee 2006; Zhou et al. 2008; Auerbach et al. 2008]. Another line of work uses ideas from behavioral economics to capture observed behavior [Rong et al. 2014].

Finally, an independent literature studies GSP and the problem of generalizing it to new contexts. The canonical analyses of GSP in a quasilinear world are [Edelman et al. 2007; Varian 2007]. Attempts to generalize GSP and consider its performance outside the standard model include [Cavallo and Wilkens 2014; Hummel and McAfee 2014; Abrams et al. 2007; Aggarwal et al. 2006].

## 2. MODEL AND PRELIMINARIES

We start with the standard model used in mechanism design, changing only the bidders’ preferences over value and money. Bidder  $i \in [n]$  has value function  $v_i(o)$  for each outcome  $o$  in a space  $\mathcal{O}$ , and a preference  $\prec_i$  over value-payment pairs  $(v_i, p_i)$  where  $p_i \in \mathbb{R}$ . The value  $v_i$  will typically be private information;  $\prec_i$  may be public or private, for example, in a quasi-linear setting,  $\prec_i$  is simply a preference for larger  $v_i - p_i$ , and is public.



Single-parameter domains are a special case of the general setting where the value function  $v_i$  is decomposed as  $v_i(o) = t_i x_i(o)$  for a private type  $t_i$  and an allocation  $x_i(o)$  associated with the outcome of the mechanism.

### Mechanisms and Truthfulness

A mechanism specifies, for bids  $b$  submitted by the bidders, an outcome  $f(b)$  and payments  $p_i(b)$ . The function  $f$  is known as the *social choice function*. We will often suppress  $f(\cdot)$  when it is clear from context, e.g. writing  $v_i(b)$  instead of  $v_i(f(b))$ .

A mechanism is *truthful* if reporting a bidder's true private information (e.g. bidding  $b_i = v_i$ ) is a dominant strategy:

**Definition 2.1** (*Dominant Strategy Incentive Compatible*). A mechanism  $\mathcal{M} = (a, p)$  is *dominant strategy incentive compatible* (DSIC) if and only if

$$(v_i(f(b_{-i}, v_i)), p_i(b_{-i}, v_i)) \succeq_i v_i(f(b_{-i}, b_i)), p_i(b_{-i}, b_i)$$

In this paper, we will have trouble in continuous bid spaces when ties occur; to solve this we introduce a slight weakening of DSIC that allows arbitrary behavior on a set of bids that “never occurs”:

**Definition 2.2** (*Dominant Strategy Incentive Compatible Almost Everywhere*). A mechanism  $\mathcal{M} = (x, p)$  is *dominant strategy incentive compatible almost everywhere* (DSIC-AE) if, for any bids  $b_{-i}$ , we have

$$(v_i(f(b_{-i}, v_i)), p_i(b_{-i}, v_i)) \succeq_i v_i(f(b_{-i}, b_i)), p_i(b_{-i}, b_i)$$

for all types  $v_i$ , except a set of  $v_i$  with measure zero.

We loosely use truthful to describe either DSIC or DSIC-AE.

**Remark 2.3.** Discretizing the bid space would eliminate the need for DSIC-AE (no type would have measure zero), but it would complicate our characterizations because it would require that we carry around  $+\epsilon$  terms that would muddy the presentation.

### Advertising Slot Auctions

When we study slot auctions for advertising, we start with standard separable click-through-rate (CTR) framework used to study sponsored search auctions. A set of  $n$  ads compete for  $m$  slots. When ad  $i$  is shown in slot  $j$ , the user clicks on it with probability

$$\Pr[\text{click on ad } i \text{ when shown in slot } j] = \alpha_j \beta_i$$

and her value is the likelihood of a click times her value for a click:

$$v_i = \Pr[\text{click on } i\text{'s ad}] \times t_i = \alpha_{j(i)} \beta_i t_i .$$

During the auction, each advertiser  $i$  submits a bid  $b_i$ , then the sponsored search platform (the auctioneer) subsequently chooses an assignment of the  $n$  advertisers to  $m \leq n$  ad slots and sets per-click prices  $p_i$ . When ad  $i$  is clicked, she pays  $p_i$ . Consistent with the literature, we assume slots are labeled such that  $\alpha_1 > \alpha_2 > \dots > \alpha_m$  (slot 1 is referred to as the top slot) and bidders are labeled so that  $\beta_1 b_1 \geq \beta_2 b_2 \geq \dots \geq \beta_n b_n$ .

When we study status-quo auctions for sponsored search, we will use the standard separable model. In this model,  $N$  ads compete for  $M$  slots. The likelihood that the user clicks on ad  $i$  when shown in slot  $j$  is

where  $\alpha_j$  is the “slot effect” and  $\beta_i$  is the “ad effect.” In this model, a bidder's private type  $t_i$  is her value for a click. It is generally assumed that  $\alpha_1 > \alpha_2 > \dots$  and that bidders are labeled such that  $\beta_1 t_1 > \beta_2 t_2 > \dots$ .

### Advertiser Preferences

Throughout the paper we will use a few different models of bidder preferences  $\prec_i$ . First, we state the standard quasi-linear preferences:

**Definition 2.4 (Quasi-linear Preferences).** Preferences are *quasi-linear* if a bidder prefers to maximize the utility function  $u_i = v_i - p_i$ , i.e.

$$(v_i, p_i) \succ_i (v'_i, p'_i) \Leftrightarrow v_i - p_i > v'_i - p'_i$$

The majority of the paper will focus on value-maximizing preferences in various forms. The general value-maximizing preferences are defined for arbitrary constraints  $C$ :

**Definition 2.5 (Value-maximizing Preferences).** Preferences are *value-maximizing* subject constraints  $C$  if they always prefer an outcome with higher value as long as none of its constraints are violated:

$$v_i > v'_i \text{ and } (v_i, p_i) \text{ satisfies } C \Rightarrow (v_i, p_i) \succeq_i (v'_i, p'_i) ,$$

and when the value is the same, a lower price is preferred:

$$p_i < p'_i \text{ and } (v_i, p_i) \text{ satisfies } C \Rightarrow (v_i, p_i) \succeq_i (v_i, p'_i) .$$

We will primarily work with two specific types of constraints:

**Definition 2.6 (Simple Value-maximizing Preferences).** A bidder has *simple value-maximizing preferences* if the bidder is value-maximizing subject to the constraint that  $p_i \leq v_i$ .

**Definition 2.7 (Value-maximizing Preferences with ROI Constraints).** A bidder has *value-maximizing preferences with an ROI constraint*  $\gamma$  if the bidder is value-maximizing with the constraint  $\frac{v_i - p_i}{p_i} \geq \gamma$ .

**Remark 2.8.** Since a value maximizer with ROI constraint  $\gamma$  is equivalent to a value maximizer with value  $v'_i = \frac{v_i}{1+\gamma}$ , we will consider mechanisms that ask for  $v'_i$  instead of  $v_i$ . This corresponds to asking the bidder how much she is willing to pay, which is natural in practice, and eliminates the need to ask the bidder for  $\gamma$ .

For some of our proofs, we will use a family of preferences that nearly interpolates between quasi-linear and simple-value-maximizing:

**Definition 2.9 ( $\alpha$ -Hybrid Preferences).** A bidder with  $\alpha$ -hybrid preferences wishes to maximize the utility  $u_i = v_i^\alpha - p_i^\alpha$ .

**Remark 2.10.** As  $\alpha \rightarrow \infty$ ,  $\alpha$ -hybrid preferences nearly converge to value maximizing preferences. The only difference occurs when  $v = p$ :  $\infty$ -hybrid preferences are indifferent between all bundles with  $v = p$ , whereas a value maximizer prefers larger  $v$ .

We also define a property that preferences may possess:

**Definition 2.11 (Super-quasi-linear).** A preference relation  $\prec_i$  is *super-quasi-linear* if an advertiser who prefers a higher-value option under quasi-linear preferences also always prefers this option under  $\prec_i$ , i.e.

$$v_i > v'_i \text{ and } v_i - p_i \geq v'_i - p'_i \Rightarrow (v_i, p_i) \succeq_i (v'_i, p'_i) .$$

### 3. MECHANISM DESIGN FOR VALUE MAXIMIZERS

Since value maximizers do not fit the quasilinear models normally used in the literature, we first must revisit fundamental results in mechanism design. We show that three of those fundamental results have natural analogs:

- In single parameter domains, any monotone allocation is incentive compatible. The prices



- Repeatedly maximizing the bidder with the highest value is the natural analog of welfare maximization.
- In unrestricted domains, any bidder-specific monotone transformations can be applied to bidders' values before maximizing "welfare." This contrasts with quasilinear settings where affine maximizers (multiplying bidders' values by weights) are the only truthful mechanisms (Roberts's Theorem [Roberts 1979]).

### 3.1. Single Parameter Mechanisms

In standard settings, it is known that any monotone allocation rule is incentive compatible [Myerson 1981; Archer and Tardos 2001]; we find that the same is true for value maximizers. This auction has very simple prices: the payment is the minimum total value ( $v = tx$ ) required to maintain the same allocation  $x$ :

**THEOREM 3.1.** *For simple value maximizers in a single parameter domain, a mechanism is truthful and individually rational if and only if it is monotone in the standard sense. Prices are computed for a bid  $b_i$  by finding the minimum (infimum) bid  $\underline{b}_i$  that achieves the same allocation  $x_i(b_{-i}, b_i)$  and setting  $p_i(b_{-i}) = \underline{b}_i x_i(b_{-i}, b_i)$ .*

*When  $\underline{b}_i$  is always the minimum, the mechanism is DSIC; otherwise, the mechanism is DSIC-AE because a bidder with type  $\underline{b}_i$  will want to lie.*

**PROOF.** Fix other bidders' bids  $b_{-i}$  and drop them for clarity. A value maximizer with single parameter type  $t_i$  will choose a bid  $b_i$  that maximizes  $x_i(b_i)$  subject to  $p_i(b_i) \leq t_i x_i(b_i)$ . Thus, for any  $t_i$ ,  $p_i(t_i)$  must be high enough that no bidder who gets a smaller allocation wants to lie, i.e.

$$\forall z \text{ where } x_i(z) < x_i(t_i) : \quad p_i(t_i) > z x_i(z) .$$

We cautiously rewrite this constraint as

$$p_i(t_i) \geq \underline{p}_i = \sup_{z | x_i(z) < x_i(t_i)} z x_i(z) .$$

Similarly, all bidders who get  $x_i(t_i)$  must pay the same price, and it must be that any bidder who gets  $x_i(t_i)$  must be willing to pay for it, so

$$p_i(t_i) \leq \bar{p}_i = \inf_{z | x_i(z) = x_i(t_i)} z x_i(z) .$$

Note that  $\underline{p}_i \geq \bar{p}_i$ , and if  $\underline{p}_i > \bar{p}_i$  for any  $t_i$ , then some bidder has an incentive to lie. Thus, it must be that  $\underline{p}_i = \bar{p}_i = p_i(t_i)$  — this is the price defined in the theorem.

Note that  $\underline{p}_i = \bar{p}_i$  for all  $t_i$  if and only if  $x_i(t_i)$  is monotone in  $t_i$ , which gives us the monotonicity requirement (intuitively, if  $x_i$  is non-monotone, a bidder can get more by bidding less, and this will clearly be a preferable outcome for a value maximizer).

We are nearly done — we have supposedly guaranteed that no bidder prefers to lie and say type  $t_i$  — but we introduced a problem when we wrote  $p_i(t_i) \geq \underline{p}_i = \sup_{z | x_i(z) < x_i(t_i)} z x_i(z)$ . This is only equivalent to the equation above it if the set of  $z$  where  $x_i(z) < x_i(t_i)$  is open on the right. If it is closed, then a bidder with type  $\hat{t}_i = \max\{z | x_i(z) < x_i(t_i)\}$  will have an incentive to bid  $\hat{t}_i + \epsilon$ . Unfortunately, this is an unavoidable; however, this can only happen at a discontinuity in  $x$ . Since the set of discontinuities must be small for a monotone function (it has measure 0), we get truthfulness in our almost-everywhere sense (DSIC-AE) if discontinuities exist.  $\square$

### 3.2. Maximizing Welfare

In standard quasilinear settings, the celebrated VCG mechanism maximizes the sum of bidders values (i.e. it chooses  $o$  maximizing  $\sum_i v_i(o)$ ) and is truthful. We show that

an analogous mechanism is truthful for value maximizers; however, instead of maximizing the  $L^1$ -norm of the value vector as in VCG, this mechanism maximizes the  $L^\infty$ -norm. Payments are an appropriate analog of the standard externality payments:

*Definition 3.2 (Informal).* The *welfare-maximizing auction for value maximizers* proceeds as follows:

- (1) Compute the outcome  $o^*$  using the following algorithm:
  - (a) Find the outcomes that maximize the value of the highest-value bidder.
  - (b) Among the outcomes from (a), find the outcomes that maximize the value of the second-highest-value bidder.
  - (c) Repeat until all bidders have been considered and call the outcome  $o^*$ ; if more than one outcome remains, pick one arbitrarily.
- (2) Payments are computed as the following “externality:” among the bidders whose value changes when  $i$  is present in the auction, identify the bidder who gets the most value if  $i$  weren’t present; charge this value to  $i$ .<sup>3</sup>

*Example 3.3.* Suppose there are three outcomes  $\{o_1, o_2, o_3\}$  and four bidders with values  $v_1 = \{3, 3, 1\}$ ,  $v_2 = \{0.5, 1, 1\}$ ,  $v_3 = \{2, 1, 0\}$ , and  $v_4 = \{0.5, 0.5, 0.5\}$ .

The welfare maximizing auction for value maximizers chooses outcome  $o_1$  as follows:

- (1) The value of the highest-value bidder is maximized by taking either outcome  $o_1$  or  $o_2$  ( $v_1 = 3$ ).
- (2) Among outcomes  $o_1$  and  $o_2$ , the value of the second-highest-value bidder is maximized by taking  $o_1$  ( $v_3 = 2$ ).

To compute prices, observe that the auction would still choose  $o_1$  if bidder 1, 2, or 4 were removed from the auction, so these bidders pay  $p_{\{1,2,4\}} = 0$ . For bidder 3, notice that the outcome would be  $o_2$  if it were removed ( $v_2 = 1$ ). Neither bidder 1 nor bidder 4 care whether the outcome is  $o_1$  or  $o_2$ , but bidder 2 gets  $v_2 = 1$  from  $o_2$  instead of  $v_2 = 0$  from  $o_1$ . The price bidder 3 pays is bidder 2’s value for  $o_2$ , i.e.  $p_3 = v_2(o_2) = 1$ .

To show why this is the natural analog of the welfare maximizing auction, we use  $\alpha$ -hybrid preferences to interpolate between quasilinear and value-maximizer preferences. These preferences will have the useful property that the natural analog of VCG maximizes the  $L^\alpha$ -norm of welfare.

Recall (Definition 2.9) that a bidder with  $\alpha$ -hybrid preferences maximizes the utility function  $u_i = v_i^\alpha - p_i^\alpha$ . Note that for  $\alpha = 1$  these are quasilinear preferences, and in the limit when  $\alpha \rightarrow \infty$  they become value-maximizing, except in the case where  $v_i = p_i$  (this event will have measure zero in our auctions).

*Definition 3.4.* The  $L^\alpha$ -welfare-maximizing auction proceeds as follows:

- (1) Choose the outcome  $o^*$  that maximizes the  $L^\alpha$ -norm of bidders’ values, i.e.

$$o^* = a(b) = \operatorname{argmax}_o \left( \sum_i b_i(o)^\alpha \right)^{\frac{1}{\alpha}}$$

- (2) Charge “externality payments”

$$p_i(b) = \left( \sum_{j \neq i} (b_j(o_{-i}^*)^\alpha - b_j(o^*)^\alpha) \right)^{\frac{1}{\alpha}}$$

<sup>3</sup>In domains with discrete typespaces, we charge a price equal to the next type above this critical value.

where  $o_{-i}^*$  is the optimal outcome with bidder  $i$  removed.

For  $\alpha = \infty$ , we choose  $o$  to maximize the limit of the  $L^\alpha$ -norm of welfare as  $\alpha \rightarrow \infty$ , and take the limit of the payment function.

It is a straightforward exercise to verify that the limiting behavior of this auction matches the informal Definition 3.2.

**THEOREM 3.5.** *For any  $\alpha$ -hybrid utilities ( $\alpha \leq \infty$ ), the  $L^\alpha$ -welfare-maximizing auction is truthful (DSIC-AE).*

**PROOF.** When  $\alpha < \infty$ , the proof follows the standard reasoning for VCG — we can write bidder  $i$ 's utility for bidding  $b_i$  when her true value is  $v_i$ , as follows:

$$u_i(v_i, b_i) = v_i(b)^\alpha - p_i(b)^\alpha = v_i(b)^\alpha + \sum_{j \neq i} b_j(b)^\alpha - \sum_{j \neq i} b_j(b_{-i}, 0)^\alpha.$$

The standard observation is that bidder  $i$  cannot affect the last term, and the first two terms are precisely the welfare objective maximized by the mechanism, so  $i$  cannot gain by lying.

For the case where  $\alpha = \infty$ , we defer proof until Theorem 3.9. In that theorem, we show that any  $L^\infty$ -affine maximizer is truthful, of which this welfare maximizing auction is a special case.<sup>4</sup>  $\square$

### 3.3. Unrestricted Domains

For quasilinear utilities, it is known that the only mechanisms implementable in the most general settings are affine maximizers [Roberts 1979]. For value maximizers, we show that more mechanisms can be implemented truthfully. However, this is a fragile fact about value-maximizers; we show that it is not generally true of  $\alpha$ -hybrid preferences for finite  $\alpha$ .

**Definition 3.6.** The  $\{\phi_i\}$  virtual welfare maximizing mechanism for value maximizers runs the welfare maximizing mechanism (Definition 3.2) on virtual values  $\phi_i(v_i)$  to get an outcome and virtual prices  $\pi_i$ ; the final prices are  $p_i = \phi_i^{-1}(\pi_i)$ .

**THEOREM 3.7.** *The virtual welfare maximizing mechanism is truthful (DSIC-AE) for value maximizers for any continuous, invertible virtual value functions  $\phi_i$ .*

**PROOF.** Observe that the welfare maximizing mechanism proceeds in rounds. In each round, the bidder who can attain the highest virtual value becomes the “winner” and we discard outcomes that do not achieve this optimal value. Let  $i_r^*$  denote the bidder who wins round  $r$  and  $\phi_r$  denote the virtual value it receives.

Assume that values are distinct so ties are impossible, i.e. for any bidders  $i$  and  $j$  and outcomes  $o_1$  and  $o_2$ ,  $\phi_i(v_i(o_1)) \neq \phi_j(v_j(o_2))$ . Note the following:

- (a) In the round where bidder  $i$  is the winner, she gets the outcome with the highest virtual value to her — which is also the outcome with the highest true value — among all remaining outcomes.
- (b) In order to win in an earlier round  $r$ , bidder  $i$  must express a value at least  $c = \phi_i^{-1}(\phi_r) > \max_o v_i(o)$  and will pay  $p_i = c$ .

(a) implies that bidder  $i$  is getting her absolute favorite outcome among the outcomes remaining when she wins, so she cannot gain by changing the outcome in that round or a later round. (b) implies that bidder  $i$  cannot gain by winning an earlier round, because she would need to pay a price greater than her maximum value for any outcome.

<sup>4</sup>We could have relied on Theorem 3.9 for the entire proof of this theorem.

We assumed no ties. Since ties happen on a portion of the value space with measure zero, we get that this auction is truthful almost everywhere.  $\square$

As noted, this suggests that more mechanisms are implementable in unrestricted domains for value maximizers than are implementable for quasilinear bidders. We show that this is a fragile result that only holds for value maximizers — for bidders with  $\alpha$ -hybrid preferences over value and money, we show that affine maximizers are the only kind of truthful mechanism for all finite  $\alpha$ .

**Definition 3.8.** A mechanism is a  $L^\alpha$ -affine-maximizer if it maximizes  $(z(o)^\alpha + \sum_i (w_i b_i(o))^\alpha)^{\frac{1}{\alpha}}$ .

**THEOREM 3.9.** For any  $\alpha < \infty$ , when  $|\mathcal{O}| \geq 3$ , bidders are  $\alpha$ -hybrid, and bidders can have any value function  $v_i : \mathcal{O} \rightarrow \mathbb{R}^+$  (the standard unrestricted setting), the only truthful mechanisms are  $L^\alpha$ -affine-maximizers.

**PROOF.** We reduce to and from quasilinear mechanisms, then rely on [Roberts 1979] to do the rest of the work. The main observation is that a mechanism that is truthful for bidders with utilities  $u_i = v_i^\alpha - p_i^\alpha$  is equivalent to a truthful mechanism for quasilinear bidders with  $u_i = \hat{v}_i - \hat{p}_i$ , where  $\hat{v}_i = v_i^\alpha$  and  $\hat{p}_i = p_i^\alpha$ .

Formally, we first show that any truthful mechanism must be an  $L^\alpha$ -affine maximizer for  $\alpha < \infty$ . Suppose we are given a truthful mechanism  $\mathcal{M}^\alpha = (f_\alpha, p_\alpha)$  for some  $\alpha$ . Observe that by truthfulness of  $\mathcal{M}_\alpha$ , it must be that bidder  $i$  prefers reporting her true value  $v_i$  to reporting another value  $b_i$ :

$$v_i(f_\alpha(b_{-i}, v_i))^\alpha - p_{\alpha,i}(b_{-i}, v_i)^\alpha \geq v_i(f_\alpha(b))^\alpha - p_{\alpha,i}(b)^\alpha .$$

Now, we construct a mechanism  $\mathcal{M}^q$  that is truthful for quasilinear bidders with value  $\hat{v}_i$  as follows: the outcome is  $f_q(\hat{v}_i) = f_\alpha(\hat{v}_i^{\frac{1}{\alpha}})$  and the payments are  $p_{q,i} = p_{\alpha,i}^\alpha$ . To see that this is truthful, observe that a quasilinear bidder with value  $\hat{v}_i$  prefers to bid  $v_i$  than  $\hat{b}_i$ :

$$\begin{aligned} u_i &= \hat{v}_i(f_q(\hat{b}_{-i}, \hat{v}_i)) - p_{q,i}(\hat{b}_{-i}, \hat{v}_i) = \left( \hat{v}_i(f_\alpha(\hat{b}_{-i}^{\frac{1}{\alpha}}, \hat{v}_i^{\frac{1}{\alpha}}))^{\frac{1}{\alpha}} \right)^\alpha - p_{\alpha,i}(\hat{b}_{-i}^{\frac{1}{\alpha}}, \hat{v}_i^{\frac{1}{\alpha}})^\alpha \\ &\geq \left( \hat{v}_i(f_\alpha(\hat{b}^{\frac{1}{\alpha}}))^{\frac{1}{\alpha}} \right)^\alpha - p_{\alpha,i}(\hat{b}^{\frac{1}{\alpha}})^\alpha \\ &= \hat{v}_i(f_q(\hat{b})) - p_{q,i}(\hat{b}) , \end{aligned}$$

where the inequality comes from truthfulness of  $\mathcal{M}^\alpha$ .

By Roberts's Theorem,  $f_q$  must be an affine maximizer, so our reduction tells us that  $f_\alpha(\hat{b}^{\frac{1}{\alpha}})$  must also be an affine maximizer, i.e.

$$f_\alpha(\hat{b}^{\frac{1}{\alpha}}) = \operatorname{argmax}_o \hat{z}(o) + \sum_i \hat{w}_i \hat{b}_i(o) .$$

By defining  $\hat{z}(o) = z(o)^\alpha$  and  $\hat{w}_i = w_i^\alpha$  we get that  $f_\alpha$  is precisely an  $L^\alpha$ -affine-maximizer:

$$f_\alpha(b) = \operatorname{argmax}_o z(o)^\alpha + \sum_i w_i^\alpha b_i(o)^\alpha = \operatorname{argmax}_o \left( z(o)^\alpha + \sum_i w_i^\alpha b_i(o)^\alpha \right)^{\frac{1}{\alpha}} .$$

To complete the case where  $\alpha < \infty$ , it is necessary to show that any  $L^\alpha$ -affine maximizer is truthful. In the quasilinear case, this is the easy direction for Roberts's theorem and is straightforward to do from first-principles — the same is true here in the manner of Theorem 3.5. Alternatively, we could use a similar reduction to prove that any  $L^\alpha$ -affine maximizer is truthful. We leave either as an exercise.  $\square$

#### 4. MECHANISM DESIGN FOR SPONSORED SEARCH AND SLOT AUCTIONS

The model we choose to represent bidder preferences fundamentally affects the prescriptions that theory makes. Modeling bidders as value maximizers turns some standard theory on its head and gives new ideas for solving some current industry auction problems. In this section we explore three applications to sponsored search:

- We show that, contrary to standard theory, the generalized second price (GSP) auction is truthful. This is an application of [Aggarwal et al. 2009].
- We use truthful auctions for value maximizers to generalize GSP to domains where it GSP currently cannot be used.
- We use preferences that are a hybrid between quasilinear and value maximizing preferences that allow a search engine to tune its pricing based on the expectations of advertisers.

##### 4.1. GSP is Truthful

The generalized second price auction (GSP) is a textbook example of the importance of modeling bidders' preferences carefully.

GSP was developed by an engineering group at Google [Varian and Harris 2014], thinking at the time that it was the appropriate way to generalize the standard second price auction to a multi-slot setting. It was quickly decided, however, that the Vickrey-Clarke-Groves auction (VCG) was actually the right way to generalize a second price auction, and much effort has been spent to rationalize and justify the continued use of GSP. However, this reaction assumes bidders have quasilinear preferences; in fact, to the degree that value maximizers are a better model of real advertisers, we show that GSP is the natural generalization of a second price auction.

The generalized second price auction is as follows:

*Definition 4.1 (GSP).* The *generalized second price* auction (GSP) proceeds as follows:

- (1) Bidders are sorted by  $\beta_i b_i$ ; the bidder ranked  $j$  gets slot  $j$ .
- (2) A bidder is charged a price-per-click equal to the minimum bid required to keep the same slot.

Aggarwal et al. [2009] study a different generalization of quasilinear preferences and remark that GSP is truthful for bidders who buy as much as they can subject to a maximum. In our setting, this is equivalent to a simple value maximizer. Alternatively, this can also be seen as a simple special case of our Theorem 3.1.

Since value maximizers with ROI constraints are equivalent to simple value maximizers with a reduced value (see Remark 2.7), we also find that GSP is truthful for value maximizers with ROI constraints.

**THEOREM 4.2 (AGGARWAL ET AL. [2009], INTERPRETED).** *GSP is truthful<sup>5</sup> for simple value maximizers and value maximizers with ROI constraints.*

##### 4.2. Generalizing GSP Beyond Slot Auctions

As noted earlier, the algorithms for optimizing ad placement in sponsored search and other online ad platforms are getting progressively more complex. As a result, GSP does not immediately apply, and a different solution is required.

The industry standard has been to use a VCG auction in these more complicated settings (e.g. Facebook or Google's Contextual Ads [Varian and Harris 2014]); however,

---

<sup>5</sup>Effectively DSIC-AE (see Definition 2.2).

in established marketplaces like sponsored search, quickly transitioning from a GSP-based system to a VCG-based system is tricky at best and perilous at worst. Advertisers are accustomed to GSP prices, so any switch will force advertisers to revise their expectations. Moreover, GSP prices are higher than their VCG counterparts, so to the degree that advertisers do not react, there may be a significant revenue loss for the advertising platform. As a result, a line of research has focused on both ways to generalize GSP (e.g. [Cavallo and Wilkens 2014; Bachrach et al. 2014]) and to transition smoothly from GSP to VCG (e.g. [Bachrach et al. 2015]).

Our work suggests a new way to approach this problem: define GSP as the welfare maximizing auction for value maximizers and run an analogous truthful auction for the new setting:

*Definition 4.3 (Generalized GSP, V1).* The generalized second price auction (GSP) is the welfare-maximizing auction for simple value maximizers.

One problem with this definition is that it changes the objective from maximizing welfare to maximizing the  $L^\infty$ -norm of bidders' values, which may be undesirable for the search engine. While this is required for truthfulness in general settings, it is not required when (as is standard in practice) the auction is a single parameter domain. Thus, we propose the following alternative definition for practically implementing a generalized GSP auction:

*Definition 4.4 (Generalized GSP, V2).* The generalized second price auction (GSP) is the auction that maximizes welfare in the standard quasilinear sense, then charges prices that are truthful for a simple value maximizer.

Both of these definitions coincide with traditional GSP in the special case of the standard model of a sponsored search auction.

#### 4.3. Hybrid Generalized GSP

While we contend that value maximization is a better model of bidder preferences than quasilinear, it also has its own problems. For example, a value maximizer will prefer 1000 clicks at \$1 each to 999 clicks for free. We hypothesize that bidders' true preferences lie somewhere between quasilinear and value-maximizing, and propose that a practical way to generalize GSP is based on a model of hybrid preferences:

*Definition 4.5 (Hybrid Generalized GSP).* The hybrid generalized second price auction (GSP) is the auction that maximizes welfare in the standard quasilinear sense, then charges prices that are truthful for a bidder with (a)  $\alpha$ -hybrid preferences, or (b) profit-maximizing preferences with a ROI constraint of  $\gamma$ .

The parameter  $\alpha$  or  $\gamma$  can be tuned to match (a) bidders' measured preferences or (b) current marketplace operating characteristics (e.g. to match current price/quantity trade-offs).

Note that for  $\alpha = 1$  or  $\gamma = 1$ , this corresponds to truthful pricing for a quasilinear bidder, while as  $\alpha \rightarrow \infty$  or as  $\gamma \rightarrow \infty$ , it corresponds to truthful pricing for a value maximizer. (See Section 5.)

### 5. MECHANISM DESIGN FOR REAL ADVERTISERS

Any ROI constraint makes a bidder look more like a simple value maximizer. For example, suppose that a bidder has an ROI constraint of 1 and has choices between an outcome  $o$  with value  $v_i(o) = 1$  and an outcome  $o'$  with value  $v_i(o') = 10$ . At any price  $p' \leq \$9$ , a profit-maximizing bidder will prefer  $o'$  to  $o$ ; however, since bidder  $i$  has an ROI constraint of 1, she will only consider outcome  $o'$  when  $p' \leq \$5$ . As a result, any



time  $o'$  is cheap enough for her to consider it, she will always prefer it to outcome  $o$  — in this example, *bidder  $i$  is effectively a simple value maximizer with value  $v'_i = \frac{1}{2}v_i$ .*

More generally, when outcomes in an auction have dramatically different values, a bidder must pay a very high price before a lesser outcome (lower value) would become preferable. In such cases, a mild ROI constraint — which caps the price an advertiser is willing to pay — will push a bidder towards simple value maximizing behavior.

The ultimate effect of ROI constraints is surprising — even if bidders are fundamentally profit-maximizers, a sufficient ROI constraint will make them value maximizers for all relevant purposes. To the degree that this is true, this again suggests that we should model advertisers as value maximizers for the purposes of designing auctions. In this section, we establish this phenomenon in theory, then show that it is very real in practice.

### 5.1. General Valuations

**THEOREM 5.1.** *When bidders have super-quasi-linear preferences  $\prec_i$ , an ROI constraint  $\gamma$ , and*

$$v_i(o) \frac{\gamma}{\gamma + 1} \geq v_i(o') \text{ for all } v_i, o, o' \text{ where } v_i(o) > v_i(o') ,$$

*no bidder wishes to lie in an auction that is truthful, individually rational for value maximizers, and has no positive transfers.<sup>6</sup>*

**PROOF.** Let  $o$  and  $p_i$  be the outcome and price that the bidder sees if it reports truthfully. Let  $o'$  and  $p'_i$  be any outcome and price that the bidder can achieve by lying.

First, suppose that  $v_i(o') < v_i(o)$ , i.e. bidder  $i$  is lying to achieve an outcome with a lower value. Then we can use individual rationality

$$p_i(b) \leq \frac{v_i(x(b))}{\gamma + 1}$$

and we know from no-positive-transfers that  $p'_i \geq 0$ . Thus, by the conditions of the theorem, we know that

$$v_i(o) - p_i \geq v_i(o) \frac{\gamma}{\gamma + 1} > v_i(o') \geq v_i(o') - p'_i .$$

In words, a profit-maximizing bidder prefers truthful reporting at the maximum individually-rational price to lying, even if  $o'$  happened for free. Since bidder  $i$  has super-profit-maximizing preferences and  $v_i(o) > v_i(o')$ , it follows  $(v_i(o), p_i) \succ_i (v_i(o'), p'_i)$ .

Next, suppose that  $v_i(o') = v_i(o)$ . We know that the mechanism is truthful for value maximizers, so it must be that  $p_i \leq p'_i$ , otherwise a value maximizer would lie to achieve  $o'$ ,  $p'_i$ . Since any super-quasi-linear bidder will prefer  $(v_i, p_i)$  to  $(v_i, p'_i)$  when  $p_i \leq p'_i$ , we can conclude that  $(v_i(o), p_i) \succeq_i (v_i(o'), p'_i)$ .

Finally, suppose that  $v_i(o') > v_i(o)$ . Since the auction is truthful for value maximizers, we know that  $p'_i > v_i(o')$ , otherwise a value maximizer would lie to achieve  $o'$ ,  $p'_i$ . By individual rationality we know  $(v_i(o), p_i) \succeq_i (0, 0)$  and any super-quasi-linear bidder prefers  $(0, 0)$  an outcome  $o'$  at price  $p' > v(o')$ . Thus,  $(v_i(o), p_i) \succeq_i (0, 0) \succ_i (v_i(o'), p'_i)$ , so bidder  $i$  will therefore prefer to tell the truth.  $\square$

<sup>6</sup>No positive transfers means that the auctioneer never pays the bidders; individual rationality says that every bidder is at least as happy as if she had not participated at all.



## 5.2. Empirical Analysis of Slot Auctions

In advertising slot auctions, a profit-maximizer’s relative value for two slots depends only on the likelihood of a click. Significantly, this gives us an easy way to test Theorem 5.1 in practice, since we can compare the likelihood of a click between slots. Formally, if we directly apply the theorem to an advertising slot auction we get the following corollary:

**COROLLARY 5.2 (OF THEOREM 5.1).** *If  $\alpha_i \geq \frac{\gamma}{\gamma+1}\alpha_{i+1}$  for all  $i$  in a standard sponsored search auction, then no bidder with super-quasi-linear preferences and a ROI constraint of  $\gamma$  wishes to lie under GSP pricing.*

This claim gives conditions under which no bidder has an incentive to lie, independent of other bidders’ bids; however, reality may be even stricter. Given bids, we can ask a simpler question of whether any bidder could possibly prefer to lie under current marketplace conditions:

**LEMMA 5.3.** *When no two bidders have the same score  $\beta t$ , and*

$$\forall i, \frac{\alpha_i}{\alpha_i - \alpha_{i+1}} - \frac{\alpha_{i+1}}{\alpha_i - \alpha_{i+1}} \frac{\beta_{i+2} b_{i+2}}{\beta_{i+1} b_{i+1}} < \gamma,$$

*no bidder with super-quasi-linear preferences and a ROI constraint of  $\gamma$  wishes to lie under GSP pricing.*

The proof (omitted) is similar to Theorem 5.1 and a version can be found in an earlier version of this work [Cavallo et al. 2015].

We tested this lemma empirically by looking at bid data for the slot auctions on Yahoo’s homepage stream. We took a dataset consisting of over one hundred thousand auctions from a brief period of time within a single day. Of course it’s impossible for us to know whether the submitted bids are reflective of true values, but we can answer the following question: assuming bidders are ROI constrained profit-maximizers, and assuming our empirically derived estimates of parameters  $\alpha$  and  $\beta$  are correct, what is the minimum value of  $\gamma$  for which Lemma 5.3 says no bidder should lie, regardless of the precise form of their preferences?

Our results are striking — if bidders require an ROI of 1, then 80% of auctions would be such that no bidder can benefit by lying under GSP pricing. This strongly suggests that GSP may in fact be the appropriate auction for this setting. See Figure 1.

## 6. CONCLUSION

The standard philosophy in theory is that no model will be perfect, but careful analysis of a simple model can be quite insightful. In mechanism design, it is typically taken for granted that a simple quasilinear model is the natural starting point for modeling any selfish agent. However, we suggest that an alternative model in which agents maximize value with minimal price sensitivity is also important. In particular, we argue that this model seems to be a substantially better fit for advertisers.

Value maximizing agents raise many problems for our current theory. Firstly, existing characterizations and existence proofs do not apply. In this paper, we largely resolved some of the fundamental questions in this vein and showed that many standard results have natural analogs. However, our focus on a single mechanism masks some deeper questions that we do not discuss:

- *How do bidders’ preferences aggregate across auctions?* With quasilinear preferences, it is reasonable to assume that a bidder’s utility is additive across different mechanisms. This is convenient because it means that an agent’s choices can be made independently. However, this assumption is problematic when bidders are




Fig. 1: Illustration of the proportion of auctions in which truth-telling is a best-response for every advertiser, given the bids of all others, assuming all bidders have a ROI constraint of  $\gamma$ . At  $\gamma = 1$ , 80% of auctions are such that nobody should lie; as  $\gamma$  approaches 2, virtually all auctions satisfy Lemma 5.3. This is derived from a dataset of auctions from Yahoo’s homepage stream, in which slot advertisements are interspersed in a stream of rich content links.

value maximizers (there is no utility function to add, for example), so it is unclear that individual mechanisms can be treated independently.

- *What are bidders’ preferences over lotteries?* Without a utility function, we cannot simply model an agent as an expected utility maximizer. Moreover, even in our hybrid preferences where a utility function exists, it is still worth asking whether expected utility maximization captures what agents seem to want.

There are also interesting open questions that build on our characterization and existence results:

- *What is the revenue-optimal auction?*
- *Does revenue equivalence hold?*

We expect that many of these questions will have answers that suggest interesting and important twists on the standard theory, just as we found that GSP is the truthful auction for value maximizers.

## REFERENCES

- ABRAMS, Z., GHOSH, A., AND VEE, E. 2007. Cost of conciseness in sponsored search auctions. In *Proceedings of the 3rd International Conference on Internet and Network Economics*. 326–334.
- ADACHI, T. 2013. Equity and the vickrey allocation rule on general preference domains. *Social Choice and Welfare* 42, 4, 813–830.

- AGGARWAL, G., GOEL, A., AND MOTWANI, R. 2006. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*. 1–7.
- AGGARWAL, G., MUTHUKRISHNAN, S., PÁL, D., AND PÁL, M. 2009. General auction mechanism for search advertising. In *Proceedings of the 18th International Conference on World Wide Web*. WWW '09. ACM, New York, NY, USA, 241–250.
- ALAEI, S., JAIN, K., AND MALEKIAN, A. 2011. Competitive equilibrium in two sided matching markets with general utility functions. *SIGecom Exch.* 10, 2, 34–36.
- ARCHER, A. AND TARDOS, E. 2001. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42Nd IEEE Symposium on Foundations of Computer Science*. FOCS '01. IEEE Computer Society, Washington, DC, USA, 482–.
- AUERBACH, J., GALENSON, J., AND SUNDARARAJAN, M. 2008. An empirical analysis of return on investment maximization in sponsored search auctions. In *Proceedings of the 2Nd International Workshop on Data Mining and Audience Intelligence for Advertising*. ADKDD '08. ACM, New York, NY, USA, 1–9.
- BACHRACH, Y., CEPPI, S., KASH, I. A., KEY, P., AND KHANI, M. R. 2015. Mechanism design for mixed ads. In *Ad Auctions Workshop*.
- BACHRACH, Y., CEPPI, S., KASH, I. A., KEY, P., AND KUROKAWA, D. 2014. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation*. 75–92.
- BORGS, C., CHAYES, J., IMMORLICA, N., JAIN, K., ETESAMI, O., AND MAHDIAN, M. 2007. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th International Conference on World Wide Web*. WWW '07. ACM, New York, NY, USA, 531–540.
- CAVALLO, R., KRISHNAMURTHY, P., AND WILKENS, C. A. 2015. On the truthfulness of gsp. In *Eleventh Workshop on Sponsored Search Auctions*.
- CAVALLO, R. AND WILKENS, C. A. 2014. *Web and Internet Economics: 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings*. Springer International Publishing, Cham, Chapter GSP with General Independent Click-through-Rates, 400–416.
- EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97, 1, 242–259.
- HUMMEL, P. AND MCAFEE, P. 2014. Position auctions with externalities and brand effects. In *Proceedings of the 10th International Workshop on Internet and Network Economics*. 585–596.
- KITTS, B. AND LEBLANC, B. 2004. Optimal bidding on keyword auctions. *Electronic Markets* 14, 3, 186–201.
- MORIMOTO, S. AND SERIZAWA, S. 2015. Strategy-proofness and efficiency with non-quasi-linear preferences: A characterization of minimum price walrasian rule. *Theoretical Economics* 10, 2, 445–487.
- MYERSON, R. B. 1981. Optimal auction design. *Math. Oper. Res.* 6, 1, 58–73.
- ROBERTS, K. 1979. The characterization of implementable social choice rules. In *Aggregation and Revelation of Preferences*, J.-J. Laffont, Ed. North-Holland, Amsterdam.
- RONG, J., QIN, T., AND AN, B. 2014. Quantal response equilibrium for sponsored search auctions: Computation and inference. In *Tenth Workshop on Sponsored Search Auctions*.
- SZYMANSKI, B. AND LEE, J. 2006. Impact of roi on bidding and revenue in sponsored search advertisement auctions. In *Second Workshop on Sponsored Search Auctions*.
- VARIAN, H. R. 2007. Position auctions. *International Journal of Industrial Organization* 25, 1163–1178.

- VARIAN, H. R. AND HARRIS, C. 2014. The vcg auction in theory and practice. *American Economic Review* 104, 5, 442–45.
- ZHOU, Y., CHAKRABARTY, D., AND LUKOSE, R. 2008. Budget constrained bidding in keyword auctions and online knapsack problems. In *Proceedings of the 17th International Conference on World Wide Web*. WWW '08. ACM, New York, NY, USA, 1243–1244.